

Robotics Research Technical Report

On Vertical Visibility in Arrangements
of Segments and the Queue Size in the
Bentley-Ottman Line Sweeping Algorithm

by

Janos Pach
Micha Sharir

Technical Report No. 417
Robotics Report No. 183
November, 1988

NYU COMPSCI TR-417

Pach, Janos
On vertical visibility in
arrangements of segments
and the queue size... c.2

New York University
Courant Institute of Mathematical Sciences

Computer Science Division
251 Mercer Street New York, N.Y. 10012



**On Vertical Visibility in Arrangements
of Segments and the Queue Size in the
Bentley-Ottman Line Sweeping Algorithm**

by

Janos Pach
Micha Sharir

Technical Report No. 417
Robotics Report No. 183
November, 1988

New York University
Dept. of Computer Science
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, New York 10012

Work on this paper has been supported by Office of Naval Research Grant N00014-87-K-0129, by National Science Foundation Grant No. NSF-DCR-83-20085, by grants from the Digital Equipment Corporation and the IBM Corporation. Work by the second author has also been supported by a grant from the U.S.-Israeli Binational Science Foundation.

On Vertical Visibility in Arrangements of Segments and the Queue Size in the Bentley-Ottman Line Sweeping Algorithm

Janos Pach^(1,2) and Micha Sharir^(1,3)

⁽¹⁾ Courant Institute of Mathematical Sciences
New York University

⁽²⁾ Mathematical Institute of the
Hungarian Academy of Sciences
and
⁽³⁾ School of Mathematical Sciences
Tel Aviv University

ABSTRACT

Let $S = \{e_1, \dots, e_n\}$ be a collection of n (intersecting) line segments in the plane. Suppose that all segments have their right endpoints lying on the same vertical line, and that we wish to bound the number of pairs of non-intersecting vertically visible segments that will intersect when extended to the right (e_i, e_j are *vertically visible* if there exists a vertical line segment connecting a point on e_i to a point on e_j and not meeting any other segment). We show that there are at most $O(n \log^2 n)$ such pairs, and only $O(n \log n)$ in the case of full rays, where the latter bound can be attained in the worst case. We apply these results to obtain similar upper and lower bounds on the maximum size of the queue in the original implementation of the Bentley-Ottman algorithm for reporting all intersections between the segments in S , i.e. the implementation where future events are not deleted from the queue. We also show that, without the extra conditions on the segments in S and on the pairs of segments to be counted, the number of non-intersecting vertically visible pairs of segments is $O(n^{4/3}(\log n)^{2/3})$, and can be $\Omega(n^{4/3})$ in the worst case.

Work on this paper has been supported by Office of Naval Research Grant N00014-87-K-0129, by National Science Foundation Grant No. NSF-DCR-83-20085, by grants from the Digital Equipment Corporation, and the IBM Corporation. Work by the second author has also been supported by a grant from the U.S.-Israeli Binational Science Foundation.

1. Introduction

Let $S = \{e_1, \dots, e_n\}$ be a collection of n line segments in the plane. The classical line-sweeping algorithm of Bentley and Ottman [BO] for reporting all k intersections of the segments in S runs in time $O((n+k) \log n)$, as follows. It maintains a priority queue Q of future events, ordered by their x coordinates, each being either an endpoint of some e_i or a detected intersection between a pair of segments in S , which occurs to the right of the (vertical) sweepline L . Each intersection event between a pair $e_i, e_j \in S$ is added to Q when e_i and e_j become adjacent along L .

(We refer to this situation by calling e_i and e_j a pair of *vertically visible* segments. Formally this means that there exists a vertical line L cutting both e_i and e_j so that the vertical segment connecting these intersections is not crossed by any other segment of S .)

In the initially proposed implementation of the algorithm, events are added to Q when the combinatorial pattern of intersections of the segments in S with L changes, which occurs when L sweeps either through an endpoint of some e_i or through an intersection of a pair e_i, e_j (in other words, when L sweeps through the currently leftmost event in Q). In each such case, only a constant number of new vertically visible pairs occur along L , and for each such pair that actually intersects to the right of L , the corresponding intersection event is added to Q . Events are removed from Q only when L sweeps through them; that is, only events at the top of Q are removed.

This strategy results in an algorithm whose running time is $O((n+k) \log n)$, where k is the total number of intersections between segments in S . The working storage of the algorithm is dominated by the maximum size of Q , which is certainly bounded by $2n+k$. Since k can be anything up to quadratic in n , this naive bound suggests the possibility that the worst-case working storage size might be as high as $\Omega(n^2)$. This has become a “folk-belief” among experts in the field, although no quadratic lower bound has ever been obtained.

To overcome this difficulty, a simple fix has been subsequently proposed by Brown [Br]. In the modified algorithm, Q contains at all times only endpoints of the segments in S ,

plus intersection events that correspond to pairs that are *currently* adjacent along L (as a matter of fact, the fix in [Br] is slightly different but achieves the same effect); this guarantees that the size of Q is always $O(n)$. This is achieved by deleting from Q every intersection event whose corresponding pair of segments are no longer vertically visible (i.e. adjacent) along L . Again, at every event swept through by L only a constant number of events have to be removed from Q , so the running time of the algorithm remains asymptotically the same. However the number of update operations on Q is essentially doubled, and the implementation of Q becomes somewhat more complicated, as we now have to provide a DELETE operation that removes elements from anywhere in the queue.

In this paper we return to the original version of the Bentley-Ottman algorithm (which does not employ the queue-deletion trick) and analyze the maximum possible size of the queue. We show that, contrary to the currently prevailing presumption, this size never exceeds $O(n \log^2 n)$. Furthermore, we show that when the algorithm is applied to a collection of lines, rather than segments, then the maximum queue size is at most $O(n \log n)$, and that this bound can be attained in the worst case. Thus, even though the size of Q can become slightly superlinear, it always remains near-linear, thus opening up the possibility of returning to the original version of the algorithm in practical applications, where the saving in the number of queue updating operations, as well as the simplicity of the data structure (which no longer requires DELETE operations to be performed for elements not in the top of the queue) may be significant.

We obtain these bounds by reducing our problem to another related one, which appears to be of independent interest, following an idea of Peter Schorn [NS]. Specifically, consider any fixed position of the sweepline L . What events are in the queue when L reaches that position? Each such event must correspond to a pair of segments that are vertically visible somewhere to the left of L and intersect to the right of L . Let us clip all segments at L , and retain only their portions to the left of L , and also discard any segment that does not reach L . Then the above observation implies that the current size of Q is bounded from above by the

number of vertically visible pairs of clipped segments of S that do not intersect one another (to the left of L , that is), but whose extensions to the right do intersect. We denote this quantity for a given collection S by $\mu(S)$. Note that in this definition all segments in S are supposed to have their right endpoints on the same vertical line (the sweepline). It is easily seen that this re-formulation of the problem involves no loss of information, in the sense that any lower bound M on $\mu(S)$ for some “vertically clipped” collection S , can be transformed into an instance of an execution of the Bentley-Ottman algorithm in which the size of Q becomes $\geq M$.

We also consider a weaker variant of the problem (which has nothing to do with the Bentley-Ottman algorithm), in which we are given an arbitrary collection of n segments and wish to estimate the number of pairs of non-intersecting vertically visible segments, dropping the condition that these pairs intersect when extended to the right (and that the segments all have to end on the same vertical line). We show that the number of these pairs in this general case is $O(n^{4/3}(\log n)^{2/3})$, and can be $\Omega(n^{4/3})$ in the worst case. (Thus the innocent-looking extra conditions that are assumed in the Bentley-Ottman case appear to be crucial for the resulting low storage bound.) This latter result is based on a random sampling technique, and its proof somewhat resembles the analysis given in [CEGSW].

The paper is organized as follows. Section 2 analyzes the case of lines, or, more generally, of a collection of segments all having the same x -projections (we refer to such configurations as *hammocks*). Section 3 analyzes the general case that arises in the Bentley-Ottman algorithm on any collection of segments, and Section 4 studies the weaker variant of vertical visibility as mentioned above. Section 5 concludes with a discussion of our results and some open problems.

2. The Case of a Hammock

Let $S = \{e_1, \dots, e_n\}$ be a collection of n segments all having the same x projection $[\xi, \eta]$. Thus their left endpoints all lie on the vertical line $L : x = \xi$, and their right endpoints

lie on the line $R : x = \eta$. Suppose the segments are sorted in increasing vertical order of their left endpoints.

(Before continuing, we note that in this case we can drop the requirement that the pairs that we wish to count intersect when extended to the right. This is because any such pair will intersect when extended either to the right or to the left (assuming no pairs of parallel segments). Thus, since the case of a hammock is symmetric with respect to the left and right directions, we can assume, without loss of generality, that at least half of the pairs we count do intersect when extended to the right.)

Define an $n \times n$ 0-1 matrix M by putting $M_{ij} = 1$ if e_i, e_j are a pair of non-intersecting vertically visible segments with e_i lying below e_j , and $M_{ij} = 0$ otherwise (in particular, M is an upper triangular matrix).

Lemma 1: M does not contain a submatrix of the form

$$\begin{bmatrix} * & 1 & 1 \\ 1 & * & 1 \end{bmatrix}$$

(where $*$ denotes any value). In other words, there do not exist two rows $a < b$ and three columns $x < y < z$ such that

$$M_{ay} = M_{az} = M_{bx} = M_{bz} = 1$$

Proof: Suppose to the contrary that M does contain such a submatrix. With a slight abuse of notation, let a, b, x, y, z also denote the corresponding segments in S . Thus $(a, y), (a, z), (b, x), (b, z)$ are all pairs of non-intersecting vertically visible segments, with a lying below y and z , and with b lying below x and z . Furthermore, denote by a_L, b_L, x_L, y_L, z_L the y coordinates of the left endpoints of these segments, and let a_R, b_R, x_R, y_R, z_R denote the y coordinates of their right endpoints. Then by definition we must have $a_L < b_L < x_L < y_L < z_L$. We next claim that a and x cannot intersect. Indeed, if they did intersect, then we would have $x_R < a_R < z_R$ (because a lies completely below z). Thus z would have to lie completely above x , which lies completely above b , so that b would not be able to see z at all, a contradiction which establishes the claim. A completely symmetric argument implies that b and y do not

intersect.

Thus the upper envelope $\psi_{a,b}$ of a and b must lie completely below the lower envelope $\phi_{x,y,z}$ of x , y and z , and any vertical visibility between a , b and x , y , z must occur between a pair of co-vertical points lying on these two respective envelopes. Consequently, each of these segments must appear along its corresponding envelope, and the vertical order of their left endpoints imply that $\psi_{a,b}$ is attained from left to right first by b and then by a , and $\phi_{x,y,z}$ is attained first by x , then by y and then by z . Let I_a , I_b , I_x , I_y , I_z denote the x -intervals where these segments appear along the corresponding envelope. Since b is assumed to see vertically both x and z , we must have $I_x \cap I_b \neq \emptyset$, $I_z \cap I_b \neq \emptyset$, which implies that $I_y \subset I_b$, which in turn contradicts the assumption that a sees y vertically, thus completing the proof of the lemma. \square

Applying the combinatorial analysis of such matrices, as given by Füredi [Fu], we obtain

Theorem 2: The maximum number of pairs of non-intersecting vertically visible segments in any collection S of n segments with the same x -projection is $\Theta(n \log n)$.

Proof: The upper bound follows immediately from Füredi's analysis. For the lower bound we use the following recursive construction. We construct collections $\{S_r\}_{r \geq 1}$ so that S_r has 2^r segments (all having $[0,1]$ as their x -projection), with $K_r \geq r \cdot 2^{r-1}$ pairs of non-intersecting vertically visible segments. S_1 is just a pair of parallel segments (with the same x -projection $[0,1]$), so $K_1 = 1$ as required. Suppose S_r has already been constructed. To obtain S_{r+1} we construct two copies of S_r . One of them, S_r^1 , is exactly S_r . The second copy S_r^2 is obtained by first rigidly translating S_r slightly upwards, and then by "shearing" it further upwards by leaving the left endpoints undisturbed and by moving each right endpoint upwards by the same very large distance c . c is chosen sufficiently large so that all intersections between segments of S_r^1 and segments of S_r^2 occur to the left of the leftmost intersection of any pair of segments in S_r . We take S_{r+1} to be $S_r^1 \cup S_r^2$. It is easily checked that for any $x \in [0,1]$ and any pair of segments e_i^1, e_j^1 in S_r^1 , if at x the segment e_i^1 lies above e_j^1 (respectively lies below e_j^1 ,

intersects e_j^1) then the same is true for the corresponding pair e_i^2, e_j^2 in S_r^2 . It follows that the number of pairs of non-intersecting vertically visible segments in S_{r+1} is at least $2K_r + 2^r$, because each $e_i^1 \in S_r^1$ and the corresponding segment $e_i^2 \in S_r^2$ form a pair of non-intersecting vertically visible segments in S_{r+1} . Thus

$$K_{r+1} \geq 2K_r + 2^r \geq (r+1) \cdot 2^r,$$

as asserted. \square

Remarks: (1) In particular, Theorem 2 implies that the queue size in the original implementation of the Bentley-Ottman algorithm, when applied to any collection of n lines, or of n segments with the same x -projection, never exceeds $O(n \log n)$.

(2) Moreover, the lower bound construction and the observation made at the beginning of this section yield an instance of the execution of the original Bentley-Ottman algorithm on a collection of n lines at which the queue size is $\Theta(n \log n)$.

3. The General Case Arising in the Bentley-Ottman Algorithm

To handle the general situation that can arise during execution of the Bentley-Ottman algorithm on an arbitrary collection of segments, we begin by considering the following special case. Suppose S and T are two collections of n segments each, such that all segments in S have a common x -projection $[\xi, \eta]$, while each segment in T has an x -projection of the form $[\zeta, \eta]$, for some $\xi < \zeta < \eta$. We refer to segments in S as "long", and to segments in T as "short". We wish to estimate the number $v(S, T)$ of pairs of non-intersecting vertically visible segments (e, e') with $e \in S, e' \in T$, with the additional requirement that e and e' would intersect when extended to the right.

Lemma 3: In the above terminology, we have $v(S, T) = \Theta(n \log n)$.

Proof: The lower bound follows immediately from Theorem 2. For the upper bound, define an $n \times n$ 0-1 matrix M as follows. Sort the segments in S in increasing vertical order of their left endpoints; let the resulting sequence be s_1, \dots, s_n . Sort the segments in T in increasing vertical order of the intersections of the lines containing them with the line $x = \xi$, and let the

resulting sequence be t_1, \dots, t_n . We now put, as before, $M_{ij} = 1$ if t_i and s_j are a pair of non-intersecting vertically visible segments, whose extensions intersect to the right of $x = \eta$, and t_i lies below s_j (a symmetric analysis will handle pairs for which t_i lies above s_j). As before, we have

Claim: M does not contain a submatrix of the form

$$\begin{bmatrix} * & 1 & 1 \\ 1 & * & 1 \end{bmatrix}$$

Indeed, suppose to the contrary that there exist segments $a, b \in T$ and $x, y, z \in S$ such that (a, y) , (a, z) , (b, x) , (b, z) are all pairs of non-intersecting vertically visible segments whose extensions intersect to the right of $x = \eta$, such that a lies below y and z , and such that b lies below x and z . Moreover, let a_R, b_R, x_R, y_R, z_R denote the y coordinates of the right endpoints of these segments, let x_L, y_L, z_L denote the y coordinates of the left endpoints of these segments, and let a_L, b_L denote the y coordinates of the intersections of the lines containing a and b with $x = \xi$. Then in the assumed configuration we have $x_L < y_L < z_L$ and $a_L < b_L$. Moreover since b and x intersect when extended to the right and b lies below x , we must also have $b_L < x_L$. Let a^* , b^* denote the extensions of a and b to the left until the line $x = \xi$ (i.e. the intersections of the lines containing a, b with the strip $\xi \leq x \leq \eta$). By assumption, a^* lies completely below y and z , and b^* lies completely below x and z (see Fig. 1). As before, we claim that a^* does not intersect x , for that would make x lying completely below z , hiding it from b^* ; similarly b^* does not intersect y . Thus any vertical visibility between a, b and x, y, z must be attained between their respective upper envelope $\psi_{a,b}$ and lower envelope $\phi_{x,y,z}$. Now $\phi_{x,y,z}$ behaves as before -- it is attained by x, y and z in this order from left to right along three respective intervals I_x, I_y, I_z . On the other hand, $\psi_{a,b}$ can now be attained by a , then b , and then a again (see Fig. 1), along three intervals I_{a1}, I_b, I_{a2} (where I_{a1} can be empty). But since b can see both x and z vertically, we must have $I_y \subset I_b$, so again it is impossible for a to see y , a contradiction which completes the proof of the claim, and thus, by [Fu], also completes the proof of the lemma. \square

Theorem 4: Let S be any collection of n line segments all having their right endpoints on the same vertical line. Then the number of pairs of non-intersecting vertically visible segments in S whose rightward extensions do intersect is $O(n \log^2 n)$.

Proof: Let $\mu(S)$ denote the number of pairs of segments in S as in the theorem statement, and let μ_n denote the maximum number of such pairs for any collection S of n segments with these properties. Assume without loss of generality that the left endpoints of the segments in S have distinct x coordinates, and let x_m denote their median value. Let S_1 be the subset of roughly $n/2$ segments whose left endpoints lie to the left of x_m , and let S_2 be the complementary subset. Then we clearly have

$$\mu(S) \leq \mu(S_1) + \mu(S_2) + \nu(S_1, S_2)$$

where $\nu(S_1, S_2)$ is the number of pairs (e_1, e_2) with $e_1 \in S_1$, and $e_2 \in S_2$ having the desired properties. By Lemma 3, this latter quantity is $O(n \log n)$, which leads to the recurrence

$$\mu_n \leq 2\mu_{n/2} + O(n \log n)$$

which solves to $\mu_n = O(n \log^2 n)$. \square

Corollary 5: The maximum queue size in the original implementation of the Bentley-Ottman algorithm, applied to any collection of n line segments, is $O(n \log^2 n)$.

Remark: We do not know whether this bound is tight in the worst case.

4. A More General Case

Although it may not be apparent from the proof of Lemma 3, it has made crucial use of the condition that the desired pairs of segments intersect when extended to the right. If we drop this condition, the number of non-intersecting vertically visible pairs can increase significantly (although still not as high as quadratic), as will be shown below.

We begin with a lower bound construction. Take an arrangement of n lines which has n faces whose total complexity is $\Theta(n^{4/3})$. Such arrangements are constructed e.g. in [EW]. For each of these faces f let $\lambda(f)$, $\rho(f)$ denote respectively the left and right portions of its boundary, delimited by the topmost and the bottommost vertices of f (see Fig. 2). Without

loss of generality we can assume that the total number of edges bounding all the left portions $\lambda(f)$ of these faces is $\Theta(n^{4/3})$. Moreover, the construction in [EW] also allows us to assume that the n faces in question are very small in size, so that they have pairwise disjoint y -projections. Next for each of these faces f draw a horizontal ray r_f extending to the right from (a point slightly to the right of) the leftmost vertex of f . Let S denote the resulting collection of $2n$ lines and rays, appropriately clipped at some vertical line sufficiently distant to the right. It is clear that for each of the special faces f and for each line l appearing along $\lambda(f)$, r_f and l are non-intersecting and vertically visible in S (the second property following from the fact that no segment r_f into another special face f'), which shows that the number of such pairs can be $\Omega(n^{4/3})$. (Note by the way that none of these pairs intersect when extended to the right.)

We next prove a closely matching upper bound, using a random sampling technique akin to that in [CEGSW]. To start the analysis we need the following variant of Lemma 3.

Lemma 6: Let S be a collection of n line segments all having x -projections contained in some interval $[\xi, \eta]$, and let $m \leq n$ be the number of “short” segments whose x -projection is not the entire $[\xi, \eta]$. Then the number of pairs of non-intersecting vertically visible segments in S is $O(mn^{1/2} + n \log n + m^{3/2}(\log m)^{1/2})$.

Proof: Let S_1 be the subset of the m short segments and S_2 the complementary subset of “long” segments. The number of desired pairs within S_2 is $O(n \log n)$ by Lemma 1. The number of such pairs (e_1, e_2) , with $e_1 \in S_1$, $e_2 \in S_2$, is analyzed as follows. Define a directed bipartite graph G between the sets S_1 , S_2 , which contains an edge (e_1, e_2) for every pair of non-intersecting vertically visible segments $e_1 \in S_1$, $e_2 \in S_2$, such that e_1 lies below e_2 . We claim that G does not contain a copy of the complete (directed) bipartite graph $K_{2,4}$ as a subgraph. Indeed, if this were the case, there would exist two short segments a , b , and four long segments e_1, \dots, e_4 such that all pairs (a, e_i) , (b, e_i) , $i = 1, \dots, 4$, have the desired properties and so that both a and b lie below all four segments e_i . Let ϕ denote the lower envelope of the four e_i ’s and let ψ denote the upper envelope of a and b . ϕ has four intervals on the x

axis so that over each of them it is attained by a fixed e_i , and ψ also has at most four such intervals so that it is attained over each of them by one of the segments a, b (see e.g. Fig. 3). By overlapping the intervals of ϕ with those of ψ and by considering all possible forms of ψ , it is easily checked that it is impossible to obtain all eight pairs of vertical visibility between a, b and the e_i 's. An application of the extremal graph-theoretic lemma of [KST] implies that the number of desired pairs (e_1, e_2) , with $e_1 \in S_1, e_2 \in S_2$, is $O(mn^{1/2} + n)$.

Finally we estimate the number of desired pairs within S_1 . Assume without loss of generality that all the endpoints of the segments in S_1 lying strictly between $x = \xi$ and $x = \eta$ have distinct x coordinates. Partition the plane into $k = (m/\log m)^{1/2}$ vertical slabs $\sigma_1, \dots, \sigma_k$ so that each of them contains at most $2(m \log m)^{1/2}$ endpoints. Consider a fixed slab σ_i , and let $p_i \leq 2(m \log m)^{1/2}$ denote the number of segments having an endpoint in σ_i and let $q_i \leq m$ denote the number of segments that cross σ_i all the way from left to right. The number of non-intersecting vertically visible pairs among the p_i short segments in σ_i is at most $O(p_i^2) = O(m \log m)$. The number of such pairs (e, e') , with e being short and e' being long in σ_i , is, by the preceding arguments, $O(p_i q_i^{1/2} + q_i) = O(m(\log m)^{1/2})$, and the number of such pairs among the q_i long segments is, by Lemma 1, $O(q_i \log q_i) = O(m \log m)$. Summing these bounds over all k slabs, we obtain that the total number of desired pairs within S_1 is $O(m^{3/2}(\log m)^{1/2})$. This completes the proof of the lemma. \square

Theorem 7: The maximum number of pairs of non-intersecting vertically visible segments in any collection of n segments in the plane is $O(n^{4/3}(\log n)^{2/3})$.

Proof: We follow the basic approach of [CEGSW], but include here, for the sake of completeness, some details of the arguments given there. Choose a random subset R of size $r = (n/\log n)^{1/3}$ of the given segments. Extend each of these segments to a full line, form the arrangement $A(R)$ of these lines, and partition its faces into $O(r^2)$ vertical trapezoidal cells, by drawing vertical segments through each intersection point until they meet another line, as in [CEGSW]. Suppose the interior of the i -th cell c_i is cut by n_i original segments and contains m_i endpoints. If we clip these segments to within c_i , and apply Lemma 6, we deduce

that the number of non-intersecting vertically visible pairs among these n_i clipped segments is

$$O(m_i n_i^{1/2} + n_i \log n_i + m_i^{3/2} (\log m_i)^{1/2}).$$

(Note that here we may have over-estimated the global count, because we may have counted pairs of non-intersecting clipped segments, for which the full segments actually intersect.) The only pairs of non-intersecting vertically visible segments that we may have missed are those with at least one of the segments in the pair belonging to R . The contribution of each cell c_i to this extra count is easily seen to be at most $2n_i + 1$, so that, summing over all cells, the number of these additional pairs is at most $O(\sum_i n_i + r^2)$.

Hence the total number of desired pairs is

$$\sum_{i=1}^{O(r^2)} O(m_i n_i^{1/2} + n_i \log n_i + m_i^{3/2} (\log m_i)^{1/2}) + O(r^2).$$

Arguing as in [CEGSW], it is easy to show that $\sum_i n_i = O(nr)$. Indeed, $\sum_i n_i = \sum_{j=1}^n l_j$, where l_j is the number of cells crossed by the j -th segment e_j . By the horizon theorem for arrangements of lines (see e.g. [Ed]) it is easily checked that each $l_j = O(r)$, establishing the claim, which implies that

$$\sum_i n_i \log n_i = O(nr \log n).$$

The probabilistic arguments in [CEGSW] imply that there exist subsets R for which

$$\sum_i m_i n_i^{1/2} = O(m(n/r)^{1/2})$$

where $m = \sum_i m_i \leq 2n$, and

$$\sum_i m_i^{3/2} (\log m_i)^{1/2} \leq (\sum_i m_i n_i^{1/2}) \cdot (\log n)^{1/2} = O(m(n/r)^{1/2} (\log n)^{1/2}).$$

Thus the total count is

$$O(m(n/r)^{1/2} (\log n)^{1/2} + nr \log n + r^2) = O(n^{4/3} (\log n)^{2/3})$$

by our choice of r . \square

5. Conclusions

In this paper we have analyzed the maximum possible size of the queue in the original version of the Bentley-Ottman line sweeping algorithm, showing that this size never exceeds $O(n \log^2 n)$ for arbitrary segments and can be at most $O(n \log n)$ in the case of lines; moreover, this latter bound can be attained in the worst case. Our solution was based on reducing the problem to a static problem analyzing the maximum number of non-intersecting vertically visible pairs of segments that do intersect when extended to the right. We have also considered a variant of this latter problem in which the “extended intersection” condition is dropped, and have shown that in this case the number of non-intersecting vertically visible pairs never exceeds $O(n^{4/3}(\log n)^{2/3})$ and can become $\Omega(n^{4/3})$ in the worst case.

The results obtained in this paper raise several open problems. One problem is whether the bound $O(n \log^2 n)$ in Theorem 4 and Corollary 5 is actually tight in the worst case, or is just an artifact of our divide-and-conquer analysis. Another problem is whether the upper bound obtained in Theorem 7 can be improved to $O(n^{4/3})$, which would then be worst-case optimal. Yet another issue is to extend our results to arrangements of more general curves. This is a natural problem since the Bentley-Ottman algorithm also applies to such curves, and it would be nice to know that the queue size cannot become too large in these more general cases as well. Concerning this problem, we note that our results (Theorems 2 and 4) apply to collections of pseudo-lines or pseudo-segments (namely when the given curves are all x -monotone, and any pair of them intersects at most once).

Finally, what are the consequences of our results to pragmatic applications of the Bentley-Ottman algorithm? Specifically, our results suggest a trade-off between the number of queue updating operations and the maximum size of the queue, and show that it is possible to save roughly half the number of updates at the cost of potentially increasing the storage for the queue by at most an $O(\log^2 n)$ factor (moreover, the implementation of the queue will be simpler, since only INSERT and DELETE-MIN operations are now required). Do these advantages justify the potentially larger storage requirements in practical executions of the algorithm?

Acknowledgements

The authors would like to thank Peter Schorn for suggesting the problem to us and for offering the basic idea of reducing the queue analysis problem to a static one involving vertical visibility (though not quite the reduction that we have used). See [NS] for more details.

REFERENCES

[BO] J. Bentley and T. Ottman, Algorithms for reporting and counting geometric intersections, *IEEE Trans. Computers*, C-28 (1979), pp. 643-647.

[Br] K.Q. Brown, Comments on "Algorithms for reporting and counting geometric intersections", *IEEE Trans. Computers*, C-30 (1981), pp. 147-148.

[CEGSW] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and surfaces, *Proc. 29th Symp. on Foundations of Computer Science*, 1988, pp. 568-579.

[Ed] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer Verlag, Heidelberg 1987.

[EW] H. Edelsbrunner and E. Welzl, On the maximal number of edges of many faces in arrangements, *J. Comb. Theory, Ser. A* 41 (1986), pp. 159-166.

[Fu] Z. Füredi, On the number of unit distances in a convex n -gon, manuscript, 1987.

[KST] T. Kővári, V.T. Sós and P. Turán, On a problem of K. Zarankiewicz, *Colloquium Math.* 3 (1954), pp. 50-57.

[NS] J. Nievergelt and P. Schorn, Geradenprobleme mit superlinearem Wachstum, *Informatik-Spektrum* 11 (4) (1988), pp. 214 - 217.




FIGURE 1




FIGURE 2




FIGURE 3

NYU COMPSCI TR-417
Pach, Janos
On vertical visibility in
arrangements of segments
and the queue size... c.2

NYU COMPSCI TR-417
Pach, Janos
On vertical visibility in
arrangements of segments
and the queue size... c.2

DATE DUE	BORROWER'S NAME

